Quantum Bootstrapping via Compressed Hamiltonian Learning

Nathan Wiebe,[1,*] Christopher Granade,[2,3,†] and David G. Cory[4,3,5,6]

[1]*Quantum Architectures and Computation Group, Microsoft Research, Redmond, WA 98052, USA*
[2]*Department of Physics, University of Waterloo, Ontario N2L 3G1, Canada*
[3]*Institute for Quantum Computing, University of Waterloo, Ontario N2L 3G1, Canada*
[4]*Department of Chemistry, University of Waterloo, Ontario N2L 3G1, Canada*
[5]*Perimeter Institute, University of Waterloo, Ontario N2L 2Y5, Canada*
[6]*Quantum Information Science Program, Canadian Institute for Advanced Research, Toronto, ON, Canada*

Recent work has shown that quantum simulation is a valuable tool for learning empirical models for quantum systems. We build upon these results by showing that a small quantum simulators can be used to characterize and learn control models for larger devices for wide classes of physically realistic Hamiltonians. Our protocol achieves this by applying Bayesian inference in concert with Lieb-Robinson bounds and interactive quantum learning methods to show that small quantum simulators can be used to efficiently compute the likelihoods needed to infer a model for a much larger quantum device. We further show that the inversion steps used in our algorithms lead to effective Lieb-Robinson velocities that are *epistemic*, in that they depend on the algorithm's uncertainty in the system Hamiltonian rather than the Hamiltonian itself. This causes the effective light cones about local observables tighten as the algorithm learns more about the system, which allows small quantum simulators to simulate their dynamics longer and in turn allows for learning to proceed at an exponential rate. We then use this approach to learn control maps and thereby provide a scalable method for characterizing and controlling large systems.

Building a large scale quantum computer or quantum simulator with existing technology seems to be a near-herculean engineering challenge. Despite these difficulties, rapid progress has been made within the last few years towards building computationally useful devices that promise to revolutionize the ways in which we solve problems in chemistry and material science, data analysis and cryptography [1–5]. Yet recently, the difficulties involved in calibrating and debugging quantum devices have suggested another possible application for a small scale quantum computer: characterizing and controlling a larger quantum computer. This process is known as quantum bootstrapping and its importance to quantum computing is perhaps best summarized by Jon Dowling [6]: "... without quantum bootstrapping it is impossible using today's classical computing resources to carefully characterize what is going on for 16 or more entangled qubits."

Here, we provide a practical quantum bootstrapping method by building on recent work showing that quantum resources, in the form of a quantum simulator or quantum computer, can be used to lead to exponential reductions in the cost of learning a Hamiltonian model for the system [7, 8], relative to state of the art methods such as classical particle filters [9, 10]. These approaches, known as Quantum Hamiltonian learning, have been shown to be robust to a wide variety of errors [8] and are surprisingly tolerant of approximation errors in the simulator. A major limitation, however, is that they seemingly require that the quantum simulator used to characterize the system is at least as large as the system of interest.

Alternative classical methods based on information locality, such as the method of Da Silva et al [11], do not suffer from this problem. However, the assumptions that enable classical learning in such methods can render them impractical for learning control maps for poorly calibrated quantum devices, which need not satisfy these assumptions. Furthermore, we show using an argument based on Fisher information that the short time evolutions used in such methods are suboptimal for characterizing devices that lack free ensemble measurements. A new method that can overcome the drawbacks of both QHL and existing classical schemes is therefore needed before quantum control and characterization of large quantum systems becomes a reality.

Our approach combines ideas from both quantum Hamiltonian learning (QHL) [7] and the information locality arguments used in [11] to circumvent the limitations of each approach alone. Our results therefore naturally lead to two distinct applications:

Compressed QHL: Learning a Hamiltonian model for a large quantum system with rapidly decaying interactions using a small quantum simulator.

Quantum bootstrapping: Designing controls for a larger quantum system with rapidly decaying interactions using a small quantum simulator.

Our work not only develops both of these applications but also provides both analytical and numerical analyses of their performance.

A phenomenon that we call *epistemic information locality* forms the basis for compressed QHL and in turn quantum bootstrapping. We make this rigorous via

* nawiebe@microsoft.com
† cgranade@cgranade.com

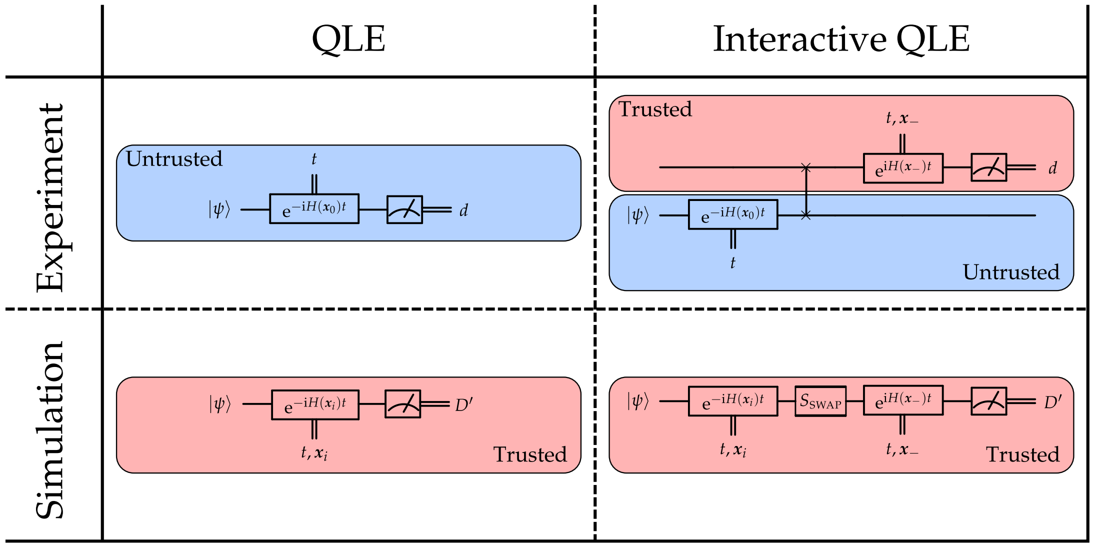

FIG. 1: Experiment and simulator design for (a) quantum Hamiltonian learning and (b) interactive quantum Hamiltonian learning with untruncated quantum simulation resources. The simulation phase is used to estimate the likelihoods of the datum found in the experiment phase.

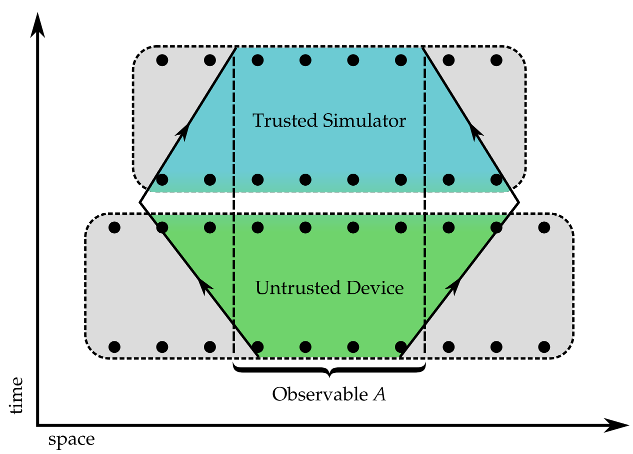

FIG. 2: Light cones for $A(t)$ for a single step of an r step protocol. The green region is the light cone after the evolution in the untrusted device, and the blue region is after inversion in the trusted device. The dashed lines show the propagation of $A(t)$ from the use of approximate inversion.

Lieb–Robinson bounds, which show that an analog of special relativity exists for local observables evolving under Hamiltonians that have rapidly decaying interactions [11–13]. Lieb–Robinson bounds give an effective "light cone", as illustrated in Figure 2, in which the evolution of an observable A can be accurately simulated without needing to consider qubits outside the light cone. Specifically, they imply that a local observable $A(t)$ provides at most an exponentially small amount of information about subsystems that are further than distance st away from the support of $A(0) \equiv A$, where s is the Lieb–Robinson velocity for the system and t is the evolution time. Here, s is analogous to the speed of light, and only depends on the geometry and strengths of the interactions in the system [12, 13]. Thus, if st is bounded above by a constant and the support of A is small then the measurement can be efficiently simulated.

We emulate long evolutions by swapping the quantum state of a subsystem of the larger (uncharacterized) system into a quantum simulator and then approximately invert the evolution using a guess for the Hamiltonian dynamics. By repeating this process rapidly, the propagation of the light cone of the observable can be delayed from extending beyond the range that can be simulated in the smaller simulator. One step of this process is illustrated in Figure 2.

In particular, we note that QHL can be used with an interactive likelihood evaluation, as shown in Figure 1, such that the observable evolves under the Hamiltonian of interest inverted by a hypothesis. We show that since the experiment design heuristic employed by QHL chooses this inversion hypothesis to be an approximation to the Hamiltonian of interest, we obtain characteristic Lieb–Robinson velocities that shrink as uncertainty about the system is reduced. *That is, the light cone represents an "epistemic" speed of light in the coupled systems that arises from the speed of information propagation depending more strongly on the uncertainty in the Hamiltonian than the Hamiltonian itself.* Since the effective speed of light slows as more information is learned, long evolutions can be used in situations where the uncertainty is small. This removes one of the main restrictions of [11].

Compressing QHL in this way is essential for our concept of quantum bootstrapping, wherein a small quantum simulator is used to *control* a larger quantum simulator, which can then be used to control an even larger simulator and so forth. Bootstrapping proceeds by us-

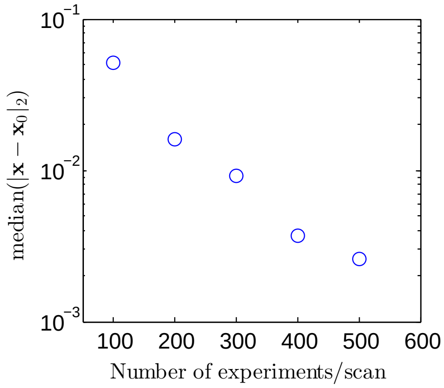

FIG. 3: Error in compressed quantum Hamiltonian learning as a function of the number of experiments per scan for 50 qubit Ising model with exponential decaying interactions.

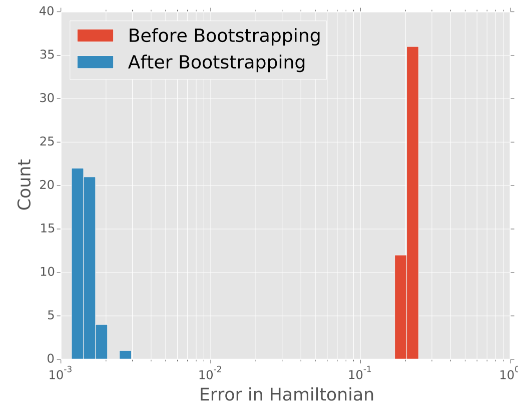

FIG. 4: Distribution of errors for each of the 49 Hamiltonian terms in the bootstrapped Hamiltonian for a 50 qubit Ising model.

ing compressed Hamiltonian learning to infer what each of the controls in the system do in isolation then uses an inversion process to find optimized controls that better implement the target dynamics. This process thereby recursively uses a small, accurate simulators to tune larger but inaccurate simulators.

In the case that the true Hamiltonian on the larger system is an affine function $H(\boldsymbol{C})$ of control settings $\boldsymbol{C}$, then we bootstrap the system by using compressed QHL to learn the Hamiltonian implemented for each of a set of different controls. That is, we start by learning the internal evolution $H(\mathbf{0})$, then proceed to learn the Hamiltonian for each of $H(\boldsymbol{C}_i)$, where $\boldsymbol{C}_i$ is taken to have support only on the i^{th} control. Then once these control maps are learned, our bootstrapping algorithm finds the settings that minimize the error in implementing a target Hamiltonian via least–squares fitting.

In order to show that our compressed quantum Hamiltonian learning and bootstrapping algorithms are scalable to large systems, we perform numerical studies with Ising-model Hamiltonians on spin chains, where interactions are taken to decay exponentially with distance. First, we apply compressed QHL to infer a Hamiltonian for unknown 50 qubit Ising-models using an 8-qubit simulator as the trusted simulator. In Figure 3, we show that the error in compressed quantum Hamiltonian learning shrinks exponentially with the number of experiments performed, in agreement with results obtained in the uncompressed case [7, 8]. Since the interactions also decay exponentially for these models, the Hamiltonian learning process is efficient.

Next, we demonstrate bootstrapping by using an 8 qubit simulator to infer the control map that describes a 50 qubit quantum simulator, from an initially uncalibrated device with crosstalk on the controls. The results shown in Figure 4 demonstrate that a modest set of experimental outcomes suffices to find controls that reduce the error in the larger device by over two orders of magnitude. These improvements are sufficiently dramatic that the bootstrapped system could realistically be used to bootstrap an even larger quantum device.

To summarize, we show that small quantum simulators can be used to characterize and calibrate larger devices, thus providing a way to bootstrap to capabilities beyond what can be implemented classically. In particular, we provide a compressed quantum Hamiltonian learning algorithm that can infer Hamiltonians for systems with local or rapidly decaying interactions. This algorithm is a necessary subroutine for bootstrapping a quantum system; wherein a small simulator to learn controls that correct Hamiltonian errors and uncertainties present in a larger quantum device. This bootstrapping protocol is useful, for instance, in calibrating control designs with cross-talk, uncertainties in coupling strengths and other effects that cause the controls on the quantum system to deviate from the designed behavior.

Our approaches, being based on quantum Hamiltonian learning, inherit the noise and sample error robustness observed in that algorithm [7, 8]. We provide numerical evidence that our techniques apply to systems with as many as 50 qubits, tolerates low precision observables and learns at an exponential rate. Thus, our quantum bootstrapping algorithm provides a scalable technique for application in even large quantum devices, and in experimentally-reasonable contexts. Our work therefore provides a critical resource for building practical quantum information processing devices and computationally useful quantum simulators.

ACKNOWLEDGMENTS

We thank Troy Borneman and Chris Ferrie for suggestions and discussions. CG was suppported by funding from Industry Canada, CERC, NSERC, CIFAR and the Province of Ontario.

[1] Nathan Wiebe, Daniel Braun, and Seth Lloyd. Quantum algorithm for data fitting. *Physical Review Letters*, 109(5):050505, 2012.

[2] Ivan Kassal, James D. Whitfield, Alejandro Perdomo-Ortiz, Man-Hong Yung, and Aln Aspuru-Guzik. Simulating chemistry using quantum computers. *Annual Review of Physical Chemistry*, 62(1):185–207, May 2011.

[3] M. B. Hastings, D. Wecker, B. Bauer, and M. Troyer. Improving quantum algorithms for quantum chemistry. *arXiv:1403.1539 [quant-ph]*, March 2014.

[4] Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM journal on computing*, 26(5):1484–1509, 1997.

[5] Brittanney Amento, Martin Rötteler, and Rainer Steinwandt. Efficient quantum circuits for binary elliptic curve arithmetic: Reducing t-gate complexity. *Quantum Info. Comput.*, 13(7-8):631–644, July 2013.

[6] Jonathan P Dowling. *Schrödinger's killer app: race to build the world's first quantum computer*. CRC Press, 2013.

[7] Nathan Wiebe, Christopher Granade, Christopher Ferrie, and D.G. Cory. Hamiltonian learning and certification using quantum resources. *Physical Review Letters*, 112(19):190501, May 2014.

[8] Nathan Wiebe, Ashish Kapoor, and Krysta Svore. Quantum nearest-neighbor algorithms for machine learning. *arXiv:1401.2142 [quant-ph]*, January 2014.

[9] Christopher E Granade, Christopher Ferrie, Nathan Wiebe, and D G Cory. Robust online hamiltonian learning. *New Journal of Physics*, 14(10):103013, October 2012.

[10] Markku P. V. Stenberg, Yuval R. Sanders, and Frank K. Wilhelm. Efficient estimation of resonant coupling between quantum systems. *arXiv:1407.5631*, July 2014.

[11] Marcus P. da Silva, Olivier Landon-Cardinal, and David Poulin. Practical characterization of quantum devices without tomography. *Physical Review Letters*, 107(21):210404, November 2011.

[12] Matthew B. Hastings and Tohru Koma. Spectral gap and exponential decay of correlations. *Communications in Mathematical Physics*, 265(3):781–804, April 2006.

[13] Bruno Nachtergaele and Robert Sims. Lieb-robinson bounds and the exponential clustering theorem. *Communications in Mathematical Physics*, 265(1):119–130, July 2006.